

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

REALTIME DATA LLC d/b/a IXO,

Plaintiff,

v.

ARRAY NETWORKS INC.,

Defendant.

C.A. No. 1:17-cv-00800-CFC

(Lead Case)

JURY TRIAL DEMANDED

REALTIME DATA LLC d/b/a IXO,

Plaintiff,

v.

MONGODB INC.,

Defendant.

C.A. No. 1:19-cv-00492-CFC

JURY TRIAL DEMANDED

**FIRST AMENDED COMPLAINT FOR PATENT INFRINGEMENT
AGAINST MONGODB, INC.**

This is an action for patent infringement arising under the Patent Laws of the United States of America, 35 U.S.C. § 1 *et seq.* in which Plaintiff Realtime Data LLC d/b/a IXO (“Plaintiff,” “Realtime,” or “IXO”) makes the following allegations against Defendant MongoDB Inc., (“MongoDB” or “Defendant”):

PARTIES

1. Realtime is a limited liability company organized under the laws of the State of New York. Realtime has a place of business at 81 Main Street, Suite 209, White Plains, NY 10601. Since the 1990s, Realtime has researched and developed specific solutions for data compression, including, for example, those that increase the speeds at which data can

be stored and accessed. As recognition of its innovations rooted in this technological field, Realtime holds a portfolio of United States patents and has numerous pending patent applications. Realtime has licensed patents in this portfolio to many of the world's leading technology companies. The patents-in-suit relate to Realtime's development of advanced systems and methods for fast and efficient data compression using numerous innovative compression techniques based on, for example, particular attributes of the data.

2. On information and belief, MongoDB is a Delaware corporation with its principal place of business at 1633 Broadway 38th Floor, New York, NY 10019. MongoDB can be served through its registered agent, Corporation Service Company, 251 Little Falls drive, Wilmington, Delaware, 19808.

JURISDICTION AND VENUE

3. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has original subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

4. This Court has personal jurisdiction over Defendant MongoDB in this action because MongoDB is incorporated in Delaware and has committed acts within the District of Delaware giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over MongoDB would not offend traditional notions of fair play and substantial justice. MongoDB, directly and through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents.

5. Venue is proper in this district under 28 U.S.C. § 1400(b). Upon information and belief, MongoDB is incorporated in Delaware, has transacted business in the District of Delaware, and has committed acts of direct and indirect infringement in this District.

ASSERTED PATENTS

6. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,116,908 (“the ’908 Patent”) entitled “System and methods for accelerated data storage and retrieval.” The ’908 Patent was duly and legally issued by the United States Patent and Trademark Office on August 25, 2015. A true and correct copy of the ’908 Patent is included as Exhibit A.

7. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,667,751 (“the ’751 Patent”) entitled “Data feed acceleration.” The ’751 Patent was duly and legally issued by the United States Patent and Trademark Office on May 30, 2017. A true and correct copy of the ’751 Patent is included as Exhibit B.

8. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein. Plaintiff Realtime is the owner by assignment of United States Patent No. 8,933,825 (“the ’825 Patent”) entitled “Data compression systems and methods.” The ’825 Patent was duly and legally issued by the United States Patent and Trademark Office on January 13, 2015. A true and correct copy of the ’825 Patent is included as Exhibit C.

9. In addition to the factual allegations set forth below for each of the four Counts, the following are non-exhaustive list of fact-based claim constructions that confirm that the claimed solutions do not just cover any form of digital data compression techniques

but instead are more focused—and covers a technical sub-species of digital data compression. These constructions include the following:¹

- a. “compressing” / “compressed” / “compression”: [representing / represented / representation of] data with fewer bits.
- b. “descriptor”: recognizable digital data
- c. “data stream”: one or more data blocks transmitted in sequence
- d. “data block”: a single unit of data, which may range in size from individual bits through complete files or collection of multiple files
- e. “analyze”: directly examine

10. Prior constructions in earlier-filed cases involving these or related patents confirm that the claimed methods and systems are in fact limited to *the compression of digital data*. For example, pursuant to a stipulation, a Texas court construed the term “compress”—a term used in all patents—to mean “represent data **with fewer bits**.” *Realtime Data LLC v. Actian Corp. et al.*, Case No. 15-cv-463-RWS-JDL, Dkt. No. 362 (E.D. Tex. July 28, 2016). This construction confirmed that the claimed inventions were limited to the realm of digital-data compression, as a “bit” is a unit of digital data. The constructions of other claim terms, such as “data block” and “accelerator” also confirmed that the patented inventions are unique to the compression of digital data. For example, the plain and ordinary meaning of the term “data block” was stipulated to be “a single **unit of data**,” which may only “range in size from individual **bits through complete files or collection of multiple files**.” *Realtime Data LLC v. Actian Corp. et al.*, Case No. 15-cv-463-RWS-JDL, Dkt. No. 362 (E.D. Tex. July 28, 2016).

11. These Patents and related patents have gone through §101 scrutiny before in multiple districts. In a detailed, twenty-two-page opinion issued on September 20, 2017,

¹ Realtime reserves the right to modify these constructions as case progresses, consistent with the practice of meeting and conferring that are typical in any claim construction proceeding.

a court in Texas ruled, in a Report and Recommendation by Magistrate Judge Love, that U.S. Patent Nos. 9,054,728, 7,415,530, and 9,116,908 are “inventive” and “directed to patent eligible subject matter” because they disclose “specific improvement[s] in computer capabilities.” *Realtime Data LLC v. Carbonite, Inc.*, Case No. 17-cv-121, D.I. 70 (E.D. Tex. Sept. 20, 2017), *e.g.*, at 7, 10, 15, 16, 20.

12. On March 7, 2018, after the Carbonite case was transferred to Massachusetts, District Judge Young in Massachusetts adopted in full Judge Love’s rulings “[a]fter careful consideration.” *Realtime Data LLC v. Carbonite, Inc.*, Case No. 1:17-cv-12499, D.I. 97 (D. Mass. March 7, 2018).

13. In addition, two judges in Texas also denied other §101 motions involving the asserted or related patents. Specifically at issue were U.S. Patent Nos. 7,378,992, 7,415,530, and 8,643,513. In one, Magistrate Judge Love held that “an assessment of the claims at issues—by a careful reading of the claims themselves—does not clearly reveal that the patents are abstract.” *Realtime Data LLC v. Actian Corp.*, 6:15-CV-463-RWS-JDL, D.I. 184 (E.D. Tex. Nov. 30, 2015). In the other, District Judge Schroeder adopted this ruling and further held that under Realtime’s view, namely, that the claims are directed to the compression of digital data, the argument that the patents are directed to an abstract idea “would fail” because the patents “provide technological solutions to problems arising specifically in the realm of computer technology.” *Realtime Data LLC v. Actian Corp.*, 6:15-CV-463-RWS-JDL, D.I. 226 (E.D. Tex. Jan. 21, 2016). Thus, in affirming the denial of the motions to dismiss, Judge Schroeder stated that, if the claim construction proceedings confirmed that the claimed inventions are specific to the methods and systems for the compression of digital data, then the claims would indeed be patent-eligible.

14. These rulings show that the patents are directed to patent eligible subject matter, and that they are also inventive.

COUNT I
INFRINGEMENT OF U.S. PATENT NO. 9,116,908

15. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein.

16. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,116,908 (“the ’908 Patent”) entitled “System and methods for accelerated data storage and retrieval.” The ’908 Patent was duly and legally issued by the United States Patent and Trademark Office on August 25, 2015. A true and correct copy of the ’908 Patent is included as Exhibit A.

17. Consistent with the claim constructions above, Plaintiff Realtime identifies the following additional constructions applicable to the claims of the ’908 patent:

- a. “data accelerator: “hardware or software with one or more compression encoders.”

18. The ’908 claims are not abstract, but rather are limited to particularized technological solutions that improve computer capabilities—e.g., digital data compression systems and methods to improve the ability of a computer system to compress and store data blocks on memory devices more efficiently. Specifically, the ’908 patent claims are directed to: computerized systems and methods for compressing and storing digital data blocks on memory devices by (1) compressing first and second data blocks with different compression techniques and (2) storing the compressed data blocks on a memory device; wherein (3) compressing and storing the compressed data blocks occurs faster than storing the uncompressed data blocks on the memory device. *See* ’908 Patent at Abstract; claim 1.

19. Because the claims are limited to the field of compressing and storing digital data blocks (each of which can comprise complete files or collection of multiple files) and are designed to improve the ability of a computer system to compress and store data blocks on memory devices more efficiently, the claims of the '908 Patent are incapable of being performed by pen and pencil. For example, it is not possible by pen and paper to compare “compression and storage time” (i.e., the time required to compress first and second data blocks with different compression techniques and storing the compressed data blocks on a memory device) against “storage time without compression” (i.e., the time required to the uncompressed data blocks on the memory device without compression). Thus, it is not possible to ensure by pen and paper that “compression and storage time” occurs faster than “storage time without compression.” Moreover, as data has been construed previously, and should be construed herein, to mean “digital data.” Accordingly, the claim herein reflect an improvement to computing technology and computers and do not reflect a previously existing solution that are simply being performed on a computer.

20. The '908 patent teaches improved, particularized digital data compression systems and methods to address problems specific to digital data. As discussed above, “data block” should be construed as “a single unit of data, which may range in size from individual bits through complete files or collection of multiple files”; and “compressing / compressed / compression” should be construed as “[representing / represented / representation] of data with fewer bits.” The '908 patent itself states that states that it deals specifically with limitations and problems arising in the realm of compressing “[d]iffuse digital data” which is “a representation of data that . . . is typically not easily recognizable to humans in its native form.” '908 patent at 1:32-36.

21. In their most basic form, and ignoring many claim limitations, the claims of the '908 patent is directed to systems of digital data compression utilizing a plurality of different compression encoders for accelerated compression and storage of data blocks. *See* '908 patent at Abstract, 2:58–3:58. The '908 patent addresses problems that existed in the realm of digital data compression, including:

- a. “high performance disk interface standards . . . offer only the promise of higher data transfer rates through intermediate data buffering in random access memory”
- b. “[f]aster disk access data rates are only achieved by the high cost solution of simultaneously accessing multiple disk drives with a technique known within the art as data striping”
- c. “problems with bandwidth limitations similarly occur within the art by all other forms of sequential, pseudorandom, and random access mass storage devices”

'908 patent at 2:19–54.

22. The '908 patent solves the foregoing problems with novel technological solutions in digital data compression utilizing a plurality of different encoders, and optionally a compression descriptor, for accelerated storage and retrieval of data blocks. The novel approaches taught in the specification, include:

- a. Using digital compression type descriptor “for output so as to indicate the type of compression format of the encoded data block”
- b. “data storage and retrieval accelerator method and system [being] employed in a disk storage adapter to reduce the time required to store and retrieve data from computer to a disk memory device”

- c. “data storage and retrieval accelerator method and system [being] employed in conjunction with random access memory to reduce the time required to store and retrieve data from random access memory”
- d. “provid[ing] an effective increase of the data storage and retrieval bandwidth of a memory storage device”

'908 patent at 2:58–3:54; 11:46–13:41. Figure 8 illustrates one preferred embodiment:

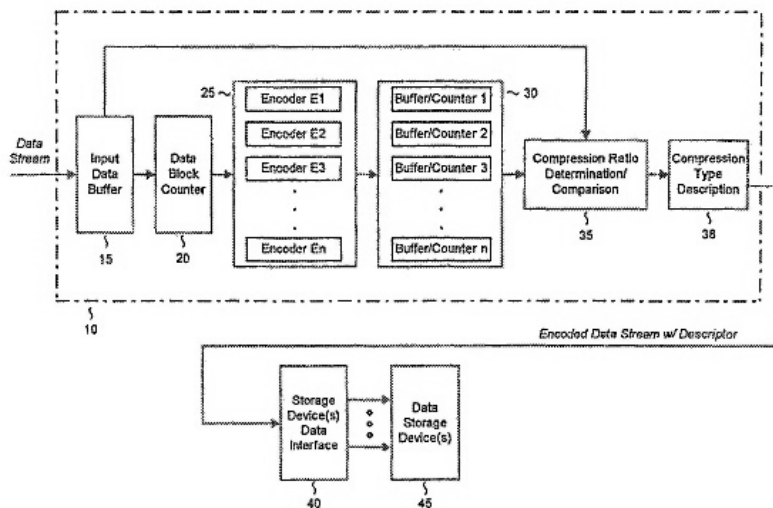


FIGURE 8

23. The claims do not merely recite a result. Instead, they recite specific steps for accomplishing a result—e.g., a data accelerator, comprising a memory device, two encoders, and a processor configured to compress two data blocks with two different compression techniques. The specification confirms this. To achieve continuous data storage acceleration, the '908 patent discloses a data storage accelerator 10 and recognizes that “the advantage of introducing a minimum delay in the time from receipt of input to storage of encoded data blocks” may be necessary in some real-time systems. '908 patent at 7:4-8:13; 7:31-34; 8:8-10. As illustrated in Figure 8, “the data storage accelerator 10 accepts data blocks from an input data stream and stores the input data block in an input

[data] buffer or cache 15.” *Id.*, 11:23-26. To compress the received data, “[t]he encoder module 25 successively receives as input each of the buffered input data blocks (or unbuffered input data blocks from the [data block counter] 20).” *Id.*, 11:52-54. “Data compression is [then] performed by the encoder module 25 wherein each of the encoders E1 . . . En processes a given input data block and outputs a corresponding set of encoded data blocks.” *Id.*, 11:54-57.

24. To achieve continuous data storage acceleration, the ’908 patent discloses several concepts that compress the received data blocks with minimal compression delay. *Id.*, 11:62-12:13; 13:22-31. First, “the encoders E1 through En of encoder module 25 may operate in parallel . . . , by executing on a plurality of processor[s] or dedicated hardware systems, or any combination thereof).” *Id.*, 11:64-12:2. Second, “a timer is included to measure the time elapsed during the encoding process against [] a priori-specified time limit” and ensure real-time or pseudo realtime encoding. *Id.*, 13:24-31. Thus, “[w]hen the time limit expires, only the data output from those encoders (in the encoder module 25) that have completed the present encoding cycle are compared to determine the encoded data with the highest compression ratio.” *Id.*, 13:26-30. Third, “encoders of the identical type may be applied in parallel to enhance encoding speed.” *Id.*, 12:9-11. “For instance, encoder E1 may comprise two parallel Huffman encoders for parallel processing of an input data block.” *Id.*, 12:11-13. By utilizing the above teachings, the data storage accelerator 10 would have allowed a system to compress and store at least two data blocks using at least two different compression techniques faster than the same system is able to store those same data blocks without any compression, as claim 1 requires.

25. In addition to continuous data storage acceleration, the '908 patent also recognizes that in some instances, the output rate of the data storage accelerator 10 may exceed the data storage rate capability of the data storage device 45. *Id.*, 9:29-10:10. To ensure compatibility in such instances, the '908 patent discloses “[t]iming and counting [the input data block to] enable[] determination of the bandwidth of the input data stream.” *Id.*, 9:27-28.) And “[i]f the bandwidths are not compatible, then one or more system parameters may be modified to make the bandwidths compatible.” *Id.*, 9:41-43. Thus a POSA would have understood from the teaching of the '908 patent's specification that (1) a compatibility problem may arise between the output rate of the data storage accelerator 10 and the write speed capability of the storage device; and (2) a solution to that problem is to make the bandwidths compatible between the data storage accelerator 10 and data storage device 45 by reducing the bandwidth of the compressed data stream so as to not exceed the bandwidth of the data storage device 45. This problem and solution is reflected in claim 26 of the '908 patent.

26. The claims require unconventional combination of elements, *e.g.*: (a) “a data accelerator” (i.e., hardware or software with one or more compression encoders) that use at least two different compression techniques; (b) “a memory device” (c) where the data accelerator is configured to compress two data blocks; (d) including “a first data block with a first compression technique”; and (e) a “second data block with a second [and different] compression technique.” The data accelerator is unconventional, as it requires compression encoders using two different compression techniques and the structural capability of compressing and storing digital data on a memory device faster than the digital data can be stored on the memory device in uncompressed form.

27. This “occurs faster than” limitation is unconventional, both alone and in an ordered combination with other claim elements. The prior art does not expressly disclose this limitation and defendants challenging the validity of the ’908 patent have not asserted that any prior art discloses this limitation or anticipates the ’908 patent. This evidenced, for example, by the outcome of IPRs on the ’908 patent and related ’530 patent. *See* IPR2016-01002 (FWD on ’908 patent), Paper 71 (PTAB Oct. 31, 2017) (affirmed by Federal Circuit); IPR2016-00972 (FWD on ’530 patent), Paper 71 (PTAB May 15, 2018) (affirmed by Federal Circuit).

28. In the FWDs, for example, the PTAB addressed obviousness challenges as to whether the combination of two prior art references rendered obvious the “occurs faster than” limitation. In those IPRs, the PTAB found that Petitioners had not demonstrated a person of ordinary skill would be motivated to combine prior art references to arrive at the occurs faster limitation. The PTAB also found that Petitioners had not demonstrated a reasonable expectation of success in the combination. As the PTAB stated: “we find that Petitioner failed to meet its burden of proof in showing: (1) how specific references could be combined, which combination(s) of elements in specific references would yield a predictable result, or how any specific combination would operate or read on the asserted claims; which is closely related to point number (2), namely, whether a person having ordinary skill in the art on this record could have had a reasonable expectation of success in attempting to combine the teachings of Franaszek with Osterlund.” IPR2016-001002, Paper 71 at 20.

29. Further, the file history confirms that the claims were inventive over prior art and not well-understood, routine, and conventional. For instance, the patent claims were

allowed by the PTO after the PTO considered hundreds of references, which are cited in the “References Cited” portion of the patent.

30. Claim 1 is not representative of all claims of the ’908 patent. In particular, the dependent claims contain limitations not found in the independent claims that further confirm that the claims are (1) are directed to digital data compression systems that improve the function and operation of a computer; and (2) recite unconventional elements taken alone or as an ordered combination. For example:

- a. Dependent claims 2–3, 22 require that the data accelerator store a first or second data descriptor is stored on the memory device indicative of the first or second compression technique, where the first / second descriptor is utilized to decompress the portion of the compressed data blocks associated with the first / second data block. Thus, claims 2–3, 22 require a data accelerator that has the structural capability of storing a data descriptor on the memory device.
- b. Dependent claims 4–5 require that the data accelerator retrieve either the first descriptor or compressed first and second compressed data blocks from the memory device. Thus, claims 4–5 requires a data accelerator that has the structural capability of retrieving a descriptor compressed data from the memory device.
- c. Dependent claim 6 requires that the data accelerator perform decompression; while dependent claims 7, 23, and 28 require that the retrieval and decompression occurs faster than retrieving the data block in compressed form.

- d. Dependent claim 8 requires that the data accelerator is coupled to the memory device via an industry standard disk interface.
 - e. Dependent claim 13 requires that the first compression technique includes compressing with a plurality of encoders “in a parallel configuration.”
 - f. Dependent claim 18 requires the first and second data blocks comprise audio or video information.
 - g. Dependent claims 19 and 25 require the first and second data blocks to be received over a communications channel.
 - h. Dependent claim 20 requires that the first compression technique is “content dependent” and the second compression technique is a form of dictionary compression.
 - i. Dependent claim 26 require determining the “bandwidth” of received data blocks; and adjusting a data rate of the compressed data blocks by “modifying a system parameter, to make a bandwidth of the compressed data blocks compatible with a bandwidth of the memory device.”
 - j. Independent claim 29 requires (1) retrieving first and second compressed data blocks compressed with different techniques from a memory device; and (2) decompressing first and second compressed data blocks; wherein (3) retrieving and decompressing the compressed data blocks occurs faster than retrieving the uncompressed data blocks from the memory device in uncompressed form.
31. In a patent filed by Altera in 2012, it admitted that there was still a technical problem associated with computer capacity and a need for a more efficient compression

system “In order to better meet the requirements of higher speed data transfer, reduced memory utilization and minimal computation in many computing applications, a need exists for computationally efficient compression and decompression.” U.S. Pat. No. 9,026,568 at 2:43-47.

32. Similarly, in a 2013 patent filed by Western Digital, it also admitted that there was still a technical problem associated with computer capacity and a need for a more efficient compression system: “It is desirable to provide mechanisms and architectures for increasing capacity, reliability, and performance of data storage systems.” U.S. Pat. No. 9,448,738 at 1:33-35.

33. The statements in these later-filed patents confirm that Realtime’s patent at issue here are directed to technical solutions to technical problems, and improves computer functionalities. The statements in these later-filed patents also confirm that the limitations recited in Realtime’s patent at issue here are not well-understood, routine, or conventional, and that the claims are not directed to other ideas “identified by the courts as abstract ideas,” that recently have been synthesized into three groups: “(a) mathematical concepts”; “(b) methods of organizing human activity”; or “(c) mental processes.” 84 Fed. Reg. 50 (Jan. 7, 2019) (2019 PTO §101 Guidance, citing and surveying post-*Alice* decisions).

34. On information and belief, MongoDB has offered for sale, sold and/or imported into the United States MongoDB products and services that infringe the ’908 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, MongoDB products and services, *e.g.*, MongoDB 3.0 and above, MongoDB Ops Manager, MongoDB Cloud Manager, WiredTiger storage

engine, and the system hardware on which they operate, and all versions and variations thereof since the issuance of the '908 Patent ("Accused Instrumentalities").

35. On information and belief, MongoDB has directly infringed and continues to infringe the '908 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities, which constitute performing a method for accelerating data storage of data claimed by Claim 21 of the '908 Patent, comprising: compressing a first data block with a first data compression technique to provide a first compressed data block; and compressing a second data block with a second data compression technique to provide a second compressed data block, wherein the first data compression technique and the second data compression technique are different; storing the first and second data compressed blocks on a memory device wherein the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form. Upon information and belief, MongoDB uses the Accused Instrumentalities, which perform the infringing method, for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

36. MongoDB also indirectly infringes the '908 Patent by manufacturing, using, selling, offering for sale, and/or importing the accused products, with knowledge that the accused products were and are especially manufactured and/or especially adapted for use in infringing the '908 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use. On information and belief, the Accused Instrumentality is designed to function with compatible hardware to perform a method for

accelerating data storage of data comprising: compressing a first data block with a first data compression technique to provide a first compressed data block; and compressing a second data block with a second data compression technique to provide a second compressed data block, wherein the first data compression technique and the second data compression technique are different; storing the first and second data compressed blocks on a memory device wherein the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form. Because the Accused Instrumentality is designed to operate as the claimed method for accelerating data storage of data, the Accused Instrumentality has no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. MongoDB's manufacture, use, sale, offering for sale, and/or importation of the Accused Instrumentality constitutes contributory infringement of the '908 Patent.

37. On information and belief, MongoDB has had knowledge of the '908 Patent since at least the filing of the original Complaint in this action, or shortly thereafter, and on information and belief, MongoDB knew of the '908 Patent and knew of its infringement, including by way of this lawsuit.

38. MongoDB's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe claims of the '908 Patent. Use of the Accused Instrumentalities in their ordinary and customary manner results in infringement of claims of the '908 Patent. For example, MongoDB explains to customers the benefits of using the Accused Instrumentalities, such as by touting their performance advantages: "[O]nly the differences

of each successive snapshot are stored. Compression and block-level deduplication reduce the size of snapshot data.” See <https://docs.opsmanager.mongodb.com/current/application/>. For similar reasons, MongoDB also induces its customers to use the Accused Instrumentalities to infringe other claims of the ’908 Patent. MongoDB specifically intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the ’908 Patent. MongoDB performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’908 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, MongoDB engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through MongoDB’s user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the ’908 Patent. Accordingly, MongoDB has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems infringing the ’908 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the ’908 Patent.

39. The Accused Instrumentalities compress a first data block with a first data compression technique to provide a first compressed data block. For example, the Accused Instrumentalities support Snappy and zlib compression techniques (*e.g.*, “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available. See <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>.) As another example, the Accused Instrumentalities provide MongoDB Ops

Manager that implements a compression technique (e.g., “Compression and block-level deduplication reduce the size of snapshot data.” *See*

<https://docs.opsmanager.mongodb.com/current/application/>.)

40. The Accused Instrumentalities compress a second data block with a second data compression technique to provide a second compressed data block, wherein the first data compression technique and the second data compression technique are different. For example, the Accused Instrumentalities include a data deduplication technique (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.)

As such, the Accused Instrumentalities define the data deduplication technique as “[A] process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.

As another example, the Accused Instrumentalities perform index prefix compression that “deduplicates common prefixes from indexed fields.” *See*

<https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>. Moreover, the Accused Instrumentalities further specify that index prefix compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

As another example, the Accused Instrumentalities support Snappy and zlib compression techniques (e.g., “By default, collection data in WiredTiger use Snappy block

compression; zlib compression is also available. *See* [https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine.\)](https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine.) As such, Snappy is a LZ77-type compressor and zlib is an LZ77 variant called deflation. *See* https://zlib.net/zlib_tech.html;https://github.com/google/snappy/blob/master/format_description.txt. In contrast to LZ77-type compression, data deduplication is a “process that eliminates redundant data.” More specifically, the Accused Instrumentalities explain that “data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>

41. The Accused Instrumentalities store the first and second data compressed blocks on a memory device wherein the compression and storage occurs faster than the first and second data blocks are able to be stored on the memory device in uncompressed form. For example, the Accused Instrumentalities store data snapshots “in a managed MongoDB database in a compressed, de-duplicated format.” *See* <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>. As another example, the Accused Instrumentalities provide routines that return the size of compressed data on disk storage (e.g., “db.collection.storageSize() to return the size in bytes of the collection on disk storage. If collection data is compressed (which is the default for WiredTiger), the storage size reflects the compressed size and may be smaller than the value returned by db.collection.dataSize().” *See* [https://docs.mongodb.com/manual/faq/storage/index.html.\)](https://docs.mongodb.com/manual/faq/storage/index.html.) As another example, the

Accused Instrumentalities state that index prefix deduplication compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” See <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>. Due to the data reduction and acceleration features of the specific compression algorithms used, the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. As such, in the Accused Instrumentalities “[W]iredTiger’s granular concurrency control and native compression will provide the best all-around performance and storage efficiency for the broadest range of applications.” See Performance Best Practices for MongoDB.

42. MongoDB also infringes other claims of the ’908 Patent, directly and through inducing infringement and contributory infringement.

43. On information and belief, use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the methods claimed by the ’908 Patent.

44. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities’ data storage accelerating features, MongoDB has injured Realtime and is liable to Realtime for infringement of the ’908 Patent pursuant to 35 U.S.C. § 271.

45. As a result of MongoDB’s infringement of the ’908 Patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for MongoDB’s infringement, but in no event less than a reasonable royalty for the use made of the invention by MongoDB, together with interest and costs as fixed by the Court.

COUNT II
INFRINGEMENT OF U.S. PATENT NO. 9,667,751

46. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein.

47. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,667,751 (“the ’751 Patent”) entitled “Data feed acceleration.” The ’751 Patent was duly and legally issued by the United States Patent and Trademark Office on May 30, 2017. A true and correct copy of the ’751 Patent is included as Exhibit B.

48. Plaintiff Realtime is the owner by assignment of United States Patent No. 9,667,751 (“the ’751 Patent”) entitled “Data feed acceleration.” The ’751 Patent was duly and legally issued by the United States Patent and Trademark Office on May 30, 2017. A true and correct copy of the ’751 Patent is included as Exhibit B.

49. The claims at issue here are not abstract, but rather are limited to particularized technological solutions that improve computer capabilities—e.g., digital data compression systems to increase the capacity of a computer system to store or transfer data more efficiently.

50. The ‘751 patent teaches various improved, particularized digital data compression systems and methods to address problems specific to digital data. Indeed, the patent itself indicate that it deals specifically with limitations and problems arising in the realm of compressing digital data. *See, e.g.*, ‘751 patent at 3:38-45. The claims are not directed merely to compressing data or merely using compression to achieve faster data storage, but also to selecting an encoder to encode data based on more than just a computer

descriptor, but instead the content of the data. This is but one of the novelties of the inventions of the '751 Patent.

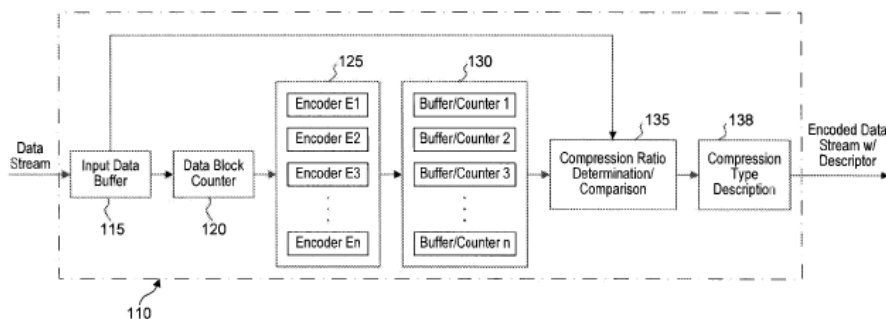
51. Because the claims are limited to the field of data blocks and are designed to increase the capacity of a computer system to store or transfer data, the claims of the '751 Patent are incapable of being performed by pen and pencil. This is reflected in the fact that data has been construed previously, and should be construed herein, to mean "digital data." Accordingly, the claim herein reflect an improvement to computing technology and computers and do not reflect a previously existing solution that are simply being performed on a computer.

52. In their most basic form, and ignoring many claim limitations, the claims of the '751 patent are directed to systems and methods for providing accelerated data transmission of digital data and effectively increasing the bandwidth of the communication channel and/or reducing the latency of data transmission. '751 patent at Abstract, 5:33–50. The '751 patent addresses specific problems in the field of optimally transmitting digital data, including:

- a. "the latency induced by the act of encryption, compression, decryption, and decompression"
- b. "substantial latency caused by aggregating data packets due to poor data compression efficiency and packet overhead"
- c. capacity limitations of data transmission using existing T1 lines
- d. "[t]he limitation of highly significant bandwidth and/or long delays with co-location processing and long latency times"

'751 patent at 1:40–5:22.

53. The '751 patent solves these and other technological problems and limitations in the prior art by providing novel technological solutions in digital data transmission, which provide, among other things, transmission and transparent multiplication of digital-data communication bandwidth, as well as a potential reduction of the latency associated with data transmission of conventional systems, and also by utilizing a state machine to compress data blocks based on an analysis of the specific content of the data being encoded. *Id.* at 5:13–29, 6:13–40. “The effective increase in bandwidth and reduction of latency of the communication channel is achieved by virtue of the faster than real-time, real-time, near real-time, compression of a received data stream prior to transmission.” *Id.* at 6:28–40. The claimed invention recognizes a characteristic, attribute, or parameter of data to select a compression encoder, and uses a state machine to provide compressed data. *Id.* Advantages of the claimed inventions include “a consistent reduction in latency” where “[t]he data compression ratio is substantial and repeatable on each data packet,” and packet independence (i.e., “no packet-to-packet data dependency”). *Id.* at 7:52–8:2. Figure 5 of the '751 patents is illustrative of one preferred embodiment:



54. To address the technological problems, the claims requires unconventional combination of elements, e.g.: (a) “identif[ying] a parameter, attribute, or value of the data

block,” (b) analysis “that excludes analyzing based solely on reading a descriptor,” (c) “selecting an encoder associated with the identified parameter, attribute, or value”; (c) “compressing data ... with the selected encoder ... utilizing a state machine”; (d) “storing compressed data block”; and (e) wherein “the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form.”

55. Further, the file history confirms that the claims were inventive over prior art and not well-understood, routine, and conventional. For instance, the patent claims were allowed by the PTO after the PTO considered hundreds of references, which are cited in the “References Cited” portion of the patent.

56. Claim 1 is not representative of all claims of the ‘751 patent. For example, claim 2 and 26 require “data packet including both control information and compressed data information,” which is not a limitation in claim 1 or other claims. As another example, claims 3 and 27 recite “utilizing Transmission Control Protocol/Internet Protocol (TCP/IP),” which is not a limitation in claim 1 or other claims. Claims 5 and 29 recite “sequence with a plurality of synchronization points.” As another example, claims 6, 7, 8, 30, 31, 32 recite “User Datagram Protocol (UDP) Packets.” Claims 10 and 33 recites “predetermined byte sequence.” Claim s11 and 34 recite “transmitting the compressed data block in a packetized data stream having data packets that include control information and compressed data information, and wherein the selected encoder is a packet independent encoder.” Claims 12, 16, 17, 18, 36, 37, 38, 39, 40, 41, 42, 44 recite “state machine” with specific configurations. As another example, claim 15 requires “transmitting the compressed data blocks in a packetized data stream of data packets having control and compressed data

information, and resetting the one or more local state machines at a predetermined point of each data packet in the packetized data stream,” which is not a limitation in claim 1 or other claims. Claims 19 and 20 recite “resetting the adaptive table at a point of each data packet in the packetized data stream.” Claims 21, 22, 45, 46 recite “lossless” encoder. Claims 23 and 47 recite “buffer.” These various limitations not recited in claim 1 confirms that claim 1 is not representative of the entire patent. Moreover, these additional limitations not recited in claim 1 further confirm that these are technological, computer improvements to technological, computer problems, as these additional limitations are not abstract but rather recite computer-based improvements to computer-based problems.

57. The claims do not merely recite a result. Instead, they recite specific steps for accomplishing a result—e.g., comprising doing analysis that excludes analyzing based solely on reading a descriptor, selecting an encoder associated with the identified parameter, attribute, or value, and utilizing a state machine, among other things. While the claims are much more than just use of a state machine, it is also noteworthy that a state machine is not an abstract idea either. For instance, state machines are specific computer solutions to computer problems. For example, Figure 3 of the ‘751 patent illustrates a “flow diagram illustrating a method for generating compression/decompression state machines according to one aspect of the present invention.” ‘751 patent at Fig. 3; 5:62-64. Figure 3 illustrates specific technological implementation, including “[p]rocess input data to acquire counts of N-tuple sequences in each global state,” “[a]pply predetermined count threshold,” and “generate a substate for each N-tuple whose count exceeds predetermined count threshold.” *Id.* The ‘751 patent further explains “using Huffman or Arithmetic encoding, wherein one or more state machine 27-27n are constructed based on a-priori knowledge of the structure

and content of one or more given broadcast and data feeds,” (‘751 patent at 9:6-10), which further confirms that state machines and other components of the ‘751 patent’s invention are technological, computer solutions to technological, computer problems, rather than merely abstract information processing.

58. The dependent claims contain limitations not found in independent claims. For example, dependent claim 2 recites “transmitting the compressed data block in a data packet to a client, the data packet including both control information and compressed data information”; claim 3 recites “wherein the compressed data block is transmitted utilizing Transmission Control Protocol/Internet Protocol (TCP/IP)”; claim 10 recites “wherein the at least one synchronization point is a predetermined byte sequence”; and claim 11 recites “transmitting the compressed data block in a packetized data stream having data packets that include control information and compressed data information, and wherein the selected encoder is a packet independent encoder.” Other claims also contain limitations not found in independent claims, as explained above.

59. In a patent filed by Altera in 2012, it admitted that there was still a technical problem associated with computer capacity and a need for a more efficient compression system: “In order to better meet the requirements of higher speed data transfer, reduced memory utilization and minimal computation in many computing applications, a need exists for computationally efficient compression and decompression.” U.S. Pat. No. 9,026,568 at 2:43-47.

60. Similarly, in a 2013 patent filed by Western Digital, it also admitted that there was still a technical problem associated with computer capacity and a need for a more efficient compression system: “It is desirable to provide mechanisms and architectures for

increasing capacity, reliability, and performance of data storage systems.” U.S. Pat. No. 9,448,738 at 1:33-35.

61. The statements in these later-filed patents confirm that Realtime’s patent at issue here are directed to technical solutions to technical problems, and improves computer functionalities. The statements in these later-filed patents also confirm that the limitations recited in Realtime’s patent at issue here are not well-understood, routine, or conventional, and that the claims are not directed to other ideas “identified by the courts as abstract ideas,” that recently have been synthesized into three groups: “(a) mathematical concepts”; “(b) methods of organizing human activity”; or “(c) mental processes.” 84 Fed. Reg. 50 (Jan. 7, 2019) (2019 PTO §101 Guidance, citing and surveying post-*Alice* decisions).

62. On information and belief, MongoDB has offered for sale, sold and/or imported into the United States MongoDB products and services that infringe the ’751 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, MongoDB products and services, *e.g.*, MongoDB 3.0 and above, MongoDB Ops Manager, MongoDB Cloud Manager, WiredTiger storage engine, and the system hardware on which they operate, and all versions and variations thereof since the issuance of the ’751 Patent (“Accused Instrumentalities”).

63. On information and belief, MongoDB has directly infringed and continues to infringe the ’751 Patent, for example, through its own use and testing of the Accused Instrumentalities, which in the ordinary course of their operation perform a method for compressing data claimed by Claim 1 of the ’751 Patent, comprising: analyzing content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor; selecting an encoder associated with the

identified parameter, attribute, or value; compressing data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine; and storing the compressed data block; wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. Upon information and belief, MongoDB uses the Accused Instrumentalities, which perform the infringing method, for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to MongoDB's customers.

64. On information and belief, MongoDB has had knowledge of the '751 Patent since at least the filing of the original Complaint in this action, or shortly thereafter, and on information and belief, MongoDB knew of the '751 Patent and knew of its infringement, including by way of this lawsuit.

65. Upon information and belief, MongoDB's affirmative acts of making, using, and selling the Accused Instrumentalities, and providing implementation services and technical support to users of the Accused Instrumentalities, have induced and continue to induce users of the Accused Instrumentalities to use them in their normal and customary way to infringe Claim 1 of the '751 Patent by analyzing content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor; selecting an encoder associated with the identified parameter, attribute, or value; compressing data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine; and storing the compressed data block; wherein the time of the compressing the data block and

the storing the compressed data block is less than the time of storing the data block in uncompressed form. For example, MongoDB explains to customers the benefits of using the Accused Instrumentalities, such as by touting their efficiency: “[O]nly the differences of each successive snapshot are stored. Compression and block-level deduplication reduce the size of snapshot data.” <https://docs.opsmanager.mongodb.com/current/application/>. For similar reasons, MongoDB also induces its customers to use the Accused Instrumentalities to infringe other claims of the ’751 Patent. MongoDB specifically intended and was aware that these normal and customary activities would infringe the ’751 Patent. MongoDB performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’751 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, MongoDB engaged in such inducement to promote the sales of the Accused Instrumentalities. Accordingly, MongoDB has induced and continues to induce users of the accused products to use the accused products in their ordinary and customary way to infringe the ’751 Patent, knowing that such use constitutes infringement of the ’751 Patent.

66. MongoDB also indirectly infringes the ’751 Patent by manufacturing, using, selling, offering for sale, and/or importing the accused products, with knowledge that the accused products were and are especially manufactured and/or especially adapted for use in infringing the ’751 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use. On information and belief, the Accused Instrumentality is designed to perform a method for compressing data comprising: analyzing content of a data block to identify a parameter, attribute, or value of the data

block that excludes analyzing based solely on reading a descriptor; selecting an encoder associated with the identified parameter, attribute, or value; compressing data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes utilizing a state machine; and storing the compressed data block; wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. Because the Accused Instrumentality is designed to operate as the claimed method for compressing, the Accused Instrumentality has no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. MongoDB's manufacture, use, sale, offering for sale, and/or importation of the Accused Instrumentality constitutes contributory infringement of the '751 Patent.

67. The Accused Instrumentalities analyze content of a data block to identify a parameter, attribute, or value of the data block that excludes analyzing based solely on reading a descriptor. For example, the Accused Instrumentalities support a data deduplication technique (e.g., "Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format." <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) As such, the Accused Instrumentalities define data deduplication as "[A] process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium." See <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>. As another example, the Accused Instrumentalities perform index prefix

compression that “deduplicates common prefixes from indexed fields.” *See* <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>. Moreover, the Accused Instrumentalities further specify that index prefix compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

68. The Accused Instrumentalities select an encoder associated with the identified parameter, attribute, or value. For example, the Accused Instrumentalities support data compression and deduplication techniques (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) As another example, the Accused Instrumentalities support index prefix deduplication compression, which “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>. As such, the Accused Instrumentalities select data deduplication technique to eliminates redundant data. In the Accused Instrumentalities, “[O]nly unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>. The Accused Instrumentalities use data compression technique to address redundancy in unique documents or blocks.

69. The Accused Instrumentalities compress data in the data block with the selected encoder to produce a compressed data block, wherein the compressing includes

utilizing a state machine. For example, the Accused Instrumentalities provide data compression and deduplication techniques outputting compressed data blocks (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.)

In particular, in the Accused Instrumentalities, deduplication is a “process that eliminates redundant data.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.

As such, the Accused Instrumentalities explain that during deduplication “[O]nly unique documents or blocks are written to a storage medium.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.

As another example, the Accused Instrumentalities support index prefix deduplication compression, which “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

As another example, the Accused Instrumentalities perform Snappy or zlib compression (e.g., “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available.” *See*

<https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>.)

70. The Accused Instrumentalities store the compressed data block. For example, the Accused Instrumentalities store data blocks “in a managed MongoDB database in a compressed, de-duplicated format.” *See*

<https://docs.opsmanager.mongodb.com/current/core/administration-interface/>. As another example, the Accused Instrumentalities provide routines that return the size of compressed

data on disk storage (e.g., “db.collection.storageSize() to return the size in bytes of the collection on disk storage. If collection data is compressed (which is the default for WiredTiger), the storage size reflects the compressed size and may be smaller than the value returned by db.collection.dataSize().” *See* <https://docs.mongodb.com/manual/faq/storage/index.html>.) As another example, the Accused Instrumentalities state that index prefix deduplication compression “[R]educes memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.

71. The Accused Instrumentalities compress data, wherein the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. For example, the Accused Instrumentalities perform compression and data de-duplication (e.g., “Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format.” <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) Due to the data reduction and acceleration features of the specific compression algorithms used, the time of the compressing the data block and the storing the compressed data block is less than the time of storing the data block in uncompressed form. As such, in the Accused Instrumentalities “[W]iredTiger’s granular concurrency control and native compression will provide the best all-around performance and storage efficiency for the broadest range of applications.” *See* Performance Best Practices for MongoDB.

72. On information and belief, MongoDB also infringes, directly and through induced infringement, and continues to infringe other claims of the ’751 Patent.

73. On information and belief, use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the methods claimed by the '751 Patent.

74. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' compression features, MongoDB has injured Realtime and is liable to Realtime for infringement of the '751 Patent pursuant to 35 U.S.C. § 271.

75. As a result of MongoDB's infringement of the '751 Patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for MongoDB's infringement, but in no event less than a reasonable royalty for the use made of the invention by MongoDB, together with interest and costs as fixed by the Court.

COUNT III
INFRINGEMENT OF U.S. PATENT NO. 8,933,825

76. Plaintiff realleges and incorporates by reference the foregoing paragraphs, as if fully set forth herein. Plaintiff Realtime is the owner by assignment of United States Patent No. 8,933,825 ("the '825 Patent") entitled "Data compression systems and methods." The '825 Patent was duly and legally issued by the United States Patent and Trademark Office on January 13, 2015. A true and correct copy of the '825 Patent is included as Exhibit C.

77. The claims at issue here are not abstract, but rather are limited to particularized technological solutions that improve computer capabilities—e.g., digital data compression systems to increase the capacity of a computer system to store or transfer data more efficiently. The claims are not directed solely to compressing data based on the

content of the data, but also to selecting an encoder to encode data based on more than just a file descriptor, but instead the content of the data. This is but one of the novelties of the inventions of the '825 Patent. This is reflected in the fact that data has been construed previously, and should be construed herein, to mean “digital data.” Accordingly, the claim herein reflect an improvement to computing technology and computers and do not reflect a previously existing solution that are simply being performed on a computer.

78. Because the claims are limited to the field of data blocks and are designed to increase the capacity of a computer system to store or transfer data, the claims of the '825 Patent are incapable of being performed by pen and pencil.

79. The '825 patent teaches various improved, particularized digital data compression systems and methods to address problems specific to digital data. Indeed, the patent itself states that it deals specifically with limitations and problems arising in the realm of compressing “[d]iffuse digital data” which is **“a representation of data that . . . is typically not easily recognizable to humans in its native form.”** '825 patent at 1:44-51.

80. In their most basic form, and ignoring many claim limitations, the claims of The '825 patent is directed to systems and methods of digital-data compression utilizing multiple encoders to compress data blocks based on an analysis of the specific content or type of the data being encoded. *See, e.g.*, '825 patent at Abstract, 3:55–5:7. The '825 patent addresses specific problems in the field of losslessly compressing digital data, including:

- a. “their content sensitive behavior . . . often referred to as data dependency”

b. “significant variations in the compression ratio obtained when using a single lossless data compression^[2] technique for data streams having different data content and data size [*i.e.*,] natural variation” ’825 patent at 1:60–3:52. The patent further explains that, while “conventional content dependent techniques may be utilized” to combat some of the problems described above, even those content dependent techniques had limitations because they relied exclusively on a descriptor such as, e.g., file extensions (e.g., “.doc,” “.txt,” etc.). The limitations included:

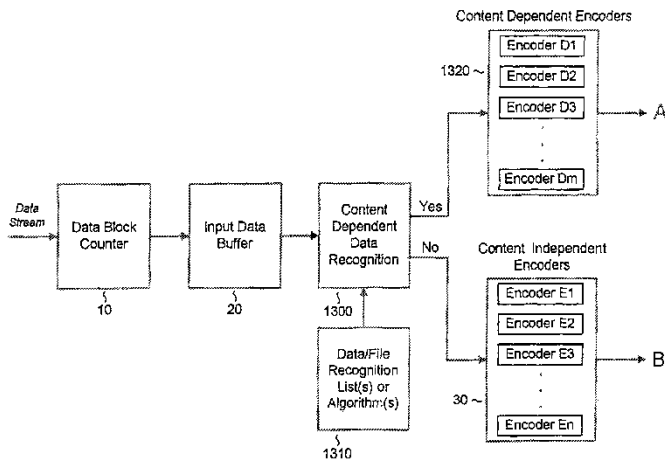
- c. “the extremely large number of application programs, some of which do not possess published or documented file formats, data structures, or data type descriptors”
- d. “the ability for any data compression supplier or consortium to acquire, store, and access the vast amounts of data required to identify known file descriptors and associated data types, data structures, and formats”
- e. “the rate at which new application programs are developed and the need to update file format data descriptions accordingly”

Id. at 3:6–19.

81. The ’825 patent solves these technological problems and others with a novel technological solution in digital-data compression utilizing multiple encoders to compress data blocks based on an analysis of the specific content or type of the data being encoded

² In “lossless” compression, “the decoded (or reconstructed) data is identical to the original uncompressed/unencoded data.” *See, e.g.*, ’825 patent at 2:18-25.

without relying solely on a descriptor such as, e.g., file extensions. For example, when one or more digital-data parameters are identified in the content of the digital data blocks, the



inventions will utilize one form of a compression encoder. And if no such digital-data parameter is identified, the inventions will utilize a different form of a compression encoder. The analysis of the digital data is not based solely a descriptor (e.g., file extensions). *See, e.g.,* ‘825 patent claim 1. Figure 13A of the ‘825 patent is illustrative of one preferred embodiment:

82. To address the technological problems, the claims requires unconventional combination of elements, e.g., (1) “wherein determining is not based solely on a descriptor that is indicative of the parameter or attribute of the data within the data block”; (2) “compressing, if the parameter or attribute of the data ... is identified, the data block with at least one encoder associated with the parameter or attribute,” and (3) “compressing, if the parameter or attribute ... is not identified, the data block with at least one encoder associated with a non-identifiable parameter or attribute.”

83. Further, the file history confirms that the claims were inventive over prior art and not well-understood, routine, and conventional. For instance, the patent claims were

allowed by the PTO after the PTO considered hundreds of references, which are cited in the “References Cited” portion of the patent.

84. Claim 1 is not representative of all claims of the ‘825 patent. For example, claim 1 requires “compressing, if the parameter or attribute of the data within the data block or attribute of the data within the data block is not identified, the data block with at least one encoder associated with a non-identifiable parameter or attribute of the data.” This is not a requirement of the other independent claims. Moreover, Claim 18 contains the different requirement of “compressing, if the first parameter or attribute of the data is not the same as one of the plurality or parameters or attributes of the data, the data block with a default encoder.” Claim 23 add the still different limitation “transmitting the compressed data block with a token indicative of the compression utilized to provide the compressed data block” which is not found in claim 1 or 18. Claim 28 provides the still different limitation “determining whether or not to compress a data block” which is not found in any of claims 1, 18, or 23.

85. The claims do not merely recite a result. Instead, they recite specific steps for accomplishing a result—e.g., comprising performing determination of parameter or attribute wherein the determining is not based solely on a descriptor that is indicative of the parameter or attribute of the data within the data block, and compressing using two different encoders based on the determined parameter or attribute, among other things.

86. The dependent claims contain limitations not found in independent claims. For example, claim 2 requires “receiving and buffering the data block, wherein the buffering is performed after the receiving of the data block and before compressing of the data block,” which is not a limitation in claim 1 or other claims. For example, dependent

claim 3 recites “transmitting a data token indicative of the compression utilized to provide the compressed data block”; claim 8 recites “wherein the size of the data block is fixed”; claim 12 recites “wherein the at least one encoder associated with the parameter or attribute of the data within the data block is lossless”; and claim 14 recites “wherein the at least one encoder associated with the parameter or attribute of the data is a Lempel-Ziv encoder.”

87. In a patent filed by Altera in 2012, it admitted that there was still a technical problem associated with computer capacity and a need for a more efficient compression system: “In order to better meet the requirements of higher speed data transfer, reduced memory utilization and minimal computation in many computing applications, a need exists for computationally efficient compression and decompression.” U.S. Pat. No. 9,026,568 at 2:43-47.

88. Similarly, in a 2013 patent filed by Western Digital, it also admitted that there was still a technical problem associated with computer capacity and a need for a more efficient compression system: “It is desirable to provide mechanisms and architectures for increasing capacity, reliability, and performance of data storage systems.” U.S. Pat. No. 9,448,738 at 1:33-35.

89. The statements in these later-filed patents confirm that Realtime’s patent at issue here are directed to technical solutions to technical problems, and improves computer functionalities. The statements in these later-filed patents also confirm that the limitations recited in Realtime’s patent at issue here are not well-understood, routine, or conventional, and that the claims are not directed to other ideas “identified by the courts as abstract ideas,” that recently have been synthesized into three groups: “(a) mathematical concepts”; “(b)

methods of organizing human activity”; or “(c) mental processes.” 84 Fed. Reg. 50 (Jan. 7, 2019) (2019 PTO §101 Guidance, citing and surveying post-*Alice* decisions).

90. On information and belief, MongoDB has offered for sale, sold and/or imported into the United States MongoDB products and services that infringe the ’825 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, MongoDB products and services, *e.g.*, MongoDB 3.0 and above, MongoDB Ops Manager, MongoDB Cloud Manager, WiredTiger storage engine, and the system hardware on which they operate, and all versions and variations thereof since the issuance of the ’825 Patent (“Accused Instrumentalities”).

91. On information and belief, MongoDB has directly infringed and continues to infringe the ’825 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities, which constitute performing a method claimed by Claim 18 of the ’825 Patent, comprising: associating at least one encoder to each one of a plurality of parameters or attributes of data; analyzing data within a data block to determine whether a parameter or attribute of the data within the data block is identified for the data block; wherein the analyzing of the data within the data block to identify a parameter or attribute of the data excludes analyzing based only on a descriptor that is indicative of the parameter or attribute of the data within the data block; identifying a first parameter or attribute of the data of the data block; compressing, if the first parameter or attribute of the data is the same as one of the plurality of parameter or attributes of the data, the data block with the at least one encoder associated with the one of the plurality of parameters or attributes of the data that is the same as the first parameter or attribute of the data to provide a compressed data

block; and compressing, if the first parameter or attribute of the data is not the same as one of the plurality of parameters or attributes of the data, the data block with a default encoder to provide the compressed data block. Upon information and belief, MongoDB uses the Accused Instrumentalities, which perform the infringing method, for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

92. MongoDB also indirectly infringes the '825 Patent by manufacturing, using, selling, offering for sale, and/or importing the accused products, with knowledge that the accused products were and are especially manufactured and/or especially adapted for use in infringing the '458 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use. On information and belief, the Accused Instrumentality is designed to function with compatible hardware to perform a method comprising: associating at least one encoder to each one of a plurality of parameters or attributes of data; analyzing data within a data block to determine whether a parameter or attribute of the data within the data block is identified for the data block; wherein the analyzing of the data within the data block to identify a parameter or attribute of the data excludes analyzing based only on a descriptor that is indicative of the parameter or attribute of the data within the data block; identifying a first parameter or attribute of the data of the data block; compressing, if the first parameter or attribute of the data is the same as one of the plurality of parameter or attributes of the data, the data block with the at least one encoder associated with the one of the plurality of parameters or attributes of the data that is the same as the first parameter or attribute of the data to provide a compressed data block; and compressing, if the first parameter or attribute of the data is not the same as one of the

plurality of parameters or attributes of the data, the data block with a default encoder to provide the compressed data block. Because the Accused Instrumentality is designed to operate as the claimed method, the Accused Instrumentality has no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. MongoDB's manufacture, use, sale, offering for sale, and/or importation of the Accused Instrumentality constitutes contributory infringement of the '825 Patent.

93. On information and belief, MongoDB has had knowledge of the '825 Patent since at least the filing of the original Complaint in this action, or shortly thereafter, and on information and belief, MongoDB knew of the '825 Patent and knew of its infringement, including by way of this lawsuit.

94. MongoDB's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe claims of the '825 Patent. Use of the Accused Instrumentalities in their ordinary and customary manner results in infringement of claims of the '825 Patent.

95. For example, MongoDB explains to customers the benefits of using the Accused Instrumentalities, such as by touting their performance advantages: "[O]nly the differences of each successive snapshot are stored. Compression and block-level deduplication reduce the size of snapshot data." <https://docs.opsmanager.mongodb.com/current/application/>. For similar reasons, MongoDB also induces its customers to use the Accused Instrumentalities to infringe other claims of the '825 Patent. MongoDB specifically intended and was aware that the normal

and customary use of the Accused Instrumentalities on compatible systems would infringe the '825 Patent. MongoDB performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '825 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, MongoDB engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through MongoDB's user manuals, product support, marketing materials, and training materials to actively induce the users of the accused products to infringe the '825 Patent. Accordingly, MongoDB has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems infringing the '825 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the '825 Patent.

96. The Accused Instrumentalities associate at least one encoder to each one of a plurality of parameters or attributes of data. For example, the Accused Instrumentalities support compression and deduplication encoders (*e.g.*, "Snapshots are stored in a managed MongoDB database in a compressed, de-duplicated format." <https://docs.opsmanager.mongodb.com/current/core/administration-interface/>.) The Accused Instrumentalities analyze data blocks to detect redundant data blocks (*e.g.*, "A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium." *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>;

“Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See*

<https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.) In particular, the Accused Instrumentalities associate deduplication encoder with the repeated data block and compression encoder with a unique data block.

97. The Accused Instrumentalities analyze data within a data block to determine whether a parameter or attribute of the data within the data block is identified for the data block. For example, the Accused Instrumentalities analyze data blocks to determine redundant data blocks by performing data deduplication. (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.) As another example, MongoDB storage engine analyzes index data structures to identify redundant data blocks (e.g., “Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.)

98. The Accused Instrumentalities analyze data, wherein the analyzing of the data within the data block to identify a parameter or attribute of the data excludes analyzing based only on a descriptor that is indicative of the parameter or attribute of the data within the data block. For example, the Accused Instrumentalities analyze data blocks to identify redundant data blocks by performing data deduplication (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or

redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>.) As another example, MongoDB storage engine analyzes index data structures to identify redundant data blocks (e.g., “Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.) As such, the Accused Instrumentalities analyze the data blocks to detect redundant data blocks.

99. The Accused Instrumentalities identify a first parameter or attribute of the data of the data block. For example, the Accused Instrumentalities analyze the data blocks to identify redundant data blocks (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See* <https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>; “Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See* <https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.)

100. The Accused Instrumentalities compress, if the first parameter or attribute of the data is the same as one of the plurality of parameter or attributes of the data, the data block with the at least one encoder associated with the one of the plurality of parameters or attributes of the data that is the same as the first parameter or attribute of the data to provide a compressed data block. For example, the Accused Instrumentalities support data

deduplication that compress data blocks by eliminates redundant data (e.g., “A process that eliminates redundant data. This data can be additional copies of database or file system documents or redundant data within those documents at the block level when written to a storage medium like spinning or solid state disks. Only unique documents or blocks are written to a storage medium.” *See*

<https://docs.opsmanager.mongodb.com/current/reference/glossary/#term-deduplication>;

“Reduces memory and disk consumption by storing any identical index key prefixes only once, per page of memory.” *See*

<https://docs.mongodb.com/manual/reference/glossary/#term-prefix-compression>.)

101. The Accused Instrumentalities compress, if the first parameter or attribute of the data is not the same as one of the plurality of parameters or attributes of the data, the data block with a default encoder to provide the compressed data block. For example, the Accused Instrumentalities provide Snappy and zlib data compression techniques (e.g., “By default, collection data in WiredTiger use Snappy block compression; zlib compression is also available. *See* <https://docs.mongodb.com/manual/faq/storage/#wiredtiger-storage-engine>.) As another example, in Accused Instrumentalities, MongoDB Ops Manager implements data compression technique during backup and archiving (e.g., “Compression and block-level deduplication reduce the size of snapshot data.” *See* <https://docs.opsmanager.mongodb.com/current/application/>.) As such, if redundant data blocks are not identified, the Accused Instrumentalities use a compression encoder to compress unique data blocks.

102. MongoDB also infringes other claims of the ’825 Patent, directly and through inducing infringement and contributory infringement.

103. On information and belief, use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the methods claimed by the '825 Patent.

104. By making, using, offering for sale, selling and/or importing into the United States the Accused Instrumentalities, and touting the benefits of using the Accused Instrumentalities' data storage accelerating features, MongoDB has injured Realtime and is liable to Realtime for infringement of the '825 Patent pursuant to 35 U.S.C. § 271.

105. As a result of MongoDB's infringement of the '825 Patent, Plaintiff Realtime is entitled to monetary damages in an amount adequate to compensate for MongoDB's infringement, but in no event less than a reasonable royalty for the use made of the invention by MongoDB, together with interest and costs as fixed by the Court.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff Realtime respectfully requests that this Court enter:

- a. A judgment in favor of Plaintiff that MongoDB has infringed, either literally and/or under the doctrine of equivalents, the '908 Patent, the '825 Patent, and the '751 Patent;
- b. A permanent injunction prohibiting MongoDB from further acts of infringement of the '908 Patent, the '825 Patent, and the '751 Patent;
- c. A judgment and order requiring MongoDB to pay Plaintiff its damages, costs, expenses, and prejudgment and post-judgment interest for its infringement of the '908 Patent, the '825 Patent, and the '751 Patent; and

d. A judgment and order requiring MongoDB to provide an accounting and to pay supplemental damages to Realtime, including without limitation, prejudgment and post-judgment interest;

e. A judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees against Defendants; and

f. Any and all other relief as the Court may deem appropriate and just under the circumstances.

DEMAND FOR JURY TRIAL

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

Dated: May 18, 2021

BAYARD, P.A.

OF COUNSEL

Marc A. Fenster (CA SBN 181067)
Reza Mirzaie (CA SBN 246953)
Paul A. Kroeger (CA SBN 229074)
C. Jay Chung (CA SBN 252794)
RUSS AUGUST & KABAT
12424 Wilshire Boulevard, 12th Floor
Los Angeles, CA 90025
(310) 826-7474
mfenster@raklaw.com
rmirzaie@raklaw.com
pkroeger@raklaw.com
jchung@raklaw.com

/s/ Stephen B. Brauerman
Stephen B. Brauerman (#4952)
Ronald P. Golden III (#6254)
600 N. King Street, Suite 400
P.O. Box 25130
Wilmington, DE 19899
Telephone: (302) 655-5000
sbrauerman@bayardlaw.com
rgolden@bayardlaw.com

*Attorneys for Plaintiff Realtime Data LLC
d/b/a IXO*

